ROBUST SATELLITE COMMUNICATIONS UNDER HOSTILE INTERFERENCE

Marc Lichtman and Jeffrey Reed

Virginia Tech
1880 Pratt Drive, Ste. 2006
Blacksburg, VA 24060

20 May 2016

Final Report

AIR FORCE RESEARCH LABORATORY
Space Vehicles Directorate
3550 Aberdeen Ave SE
AIR FORCE MATERIEL COMMAND
KIRTLAND AIR FORCE BASE, NM 87117-5776

DTIC COPY
NOTICE AND SIGNATURE PAGE

AFRL-RV-PS-TR-2016-0079 HAS BEEN REVIEWED AND IS APPROVED FOR PUBLICATION IN ACCORDANCE WITH ASSIGNED DISTRIBUTION STATEMENT.

//SIGNED//   //SIGNED//
KHANH PHAM   PAUL D. LEVAN, Ph.D.
Program Manager   Technical Advisor, Space Based Advanced Sensing
                                                             and Protection

//SIGNED//
JOHN BEAUCHEMIN
Chief Engineer, Spacecraft Technology Division
Space Vehicles Directorate

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

1. REPORT DATE *(DD-MM-YYYY)*	2. REPORT TYPE	3. DATES COVERED *(From - To)*
20-05-2016	Final Report	13 May 2014 to 30 Mar 2016

4. TITLE AND SUBTITLE	5a. CONTRACT NUMBER
Robust Satellite Communications Under Hostile Interference	FA9453-14-1-0222
	5b. GRANT NUMBER
	5c. PROGRAM ELEMENT NUMBER 62601F

6. AUTHOR(S)	5d. PROJECT NUMBER 8809
Marc Lichtman and Jeffrey Reed	**5e. TASK NUMBER** PPM00019612
	5f. WORK UNIT NUMBER EF122073

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)	8. PERFORMING ORGANIZATION REPORT NUMBER
Virginia Tech 1880 Pratt Drive, Ste. 2006 Blacksburg, VA 24060	

9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES)	10. SPONSOR/MONITOR'S ACRONYM(S)
Air Force Research Laboratory Space Vehicles Directorate 3550 Aberdeen Ave SE Kirtland AFB, NM 87117-5776	AFRL/RVSW
	11. SPONSOR/MONITOR'S REPORT NUMBER(S) AFRL-RV-PS-TR-2016-0079

12. DISTRIBUTION / AVAILABILITY STATEMENT
Approved for public release; distribution is unlimited.

13. SUPPLEMENTARY NOTES

14. ABSTRACT
Antifragility refers to systems that increase in capability, resilience, or robustness as a result of mistakes, faults, attacks, or failures. An antifragile system is fundamentally different from one that is resilient (able to recover from failure) and robust (able to resist failure). We apply the concept of antifragility to wireless communications where the system stressor is a jammer that intends to disrupt the underlying communications. Through a novel strategy, we exploit the communications jammer, providing an increase in communications capability such as bits per second. We show that an antifragile gain is possible under a wide variety of reactive-jamming scenarios and provide guidelines for realizing these gains by creating an antifragile waveform. The material in this report has been published in the Institute of Electrical and Electronics Enginers (IEEE) Systems Journal.

15. SUBJECT TERMS
reactive jamming, antifragile communications, jamming taxonomy, jammer exploitation

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON Khanh Pham
a. REPORT Unclassified	**b. ABSTRACT** Unclassified	**c. THIS PAGE** Unclassified	Unlimited	36	**19b. TELEPHONE NUMBER** *(include area code)*

Standard Form 298 (Rev. 8-98)

(This page intentionally left blank)

Table of Contents

List of Figures

ACKNOWLEDGMENTS

DISCLAIMER

The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of Air Force Research Laboratory or the U.S. Government.

(This page intentionally left blank)

1 Summary

In this report we introduce the concept of antifragile communications, which we define as the capability for a communications system to improve in performance due to a system stressor or harsh condition. Antifragility refers to systems that increase in capability, resilience, or robustness as a result of mistakes, faults, attacks, or failures. An antifragile system is fundamentally different from one that is resilient (able to recover from failure) and robust (able to resist failure). We apply the concept of antifragility to wireless communications where the system stressor is a jammer that intends to disrupt the underlying communications. Through a novel strategy, we exploit the communications jammer, providing an increase in communications capability such as bits per second. We show that an antifragile gain is possible under a wide variety of reactive-jamming scenarios and provide guidelines for realizing these gains by creating an antifragile waveform. The material in this report has been published in Institute of Electrical and Electronics Enginers (*IEEE) Systems Journal* [1].

2 Introduction

Jamming is an ongoing threat that plagues wireless communications in contested areas. Unfortunately, jamming complexity and sophistication will continue to increase over time. This is in part due to the availability of powerful and low-cost software-defined radios. The current approach to countering such a threat revolves around *jammer detection* and *jammer mitigation*. As such, an increase in jammer complexity requires an increase in countermeasure complexity. This leads to extremely hardened radios that sacrifice communications performance for more advanced jamming protection. To provide an escape from this trend, we investigate the previously unexplored area of *jammer exploitation*. Unlike mitigation (i.e., anti-jamming), the more complex an enemy jammer, the more potential there is for exploitation. It is for this reason that the antifragile paradigm should be applied to wireless communications. An example of jammer exploitation includes manipulating a jammer into jamming a particular sequence of channels, where data is conveyed in the sequence of channels selected, similar to frequency shift keying. A strategy that exploits a jamming attack to provide a communications gain, such as reducing the bit error rate or increasing the rate of communication, can be labeled as "antifragile communications". This is because a gain is achieved by harnessing the presence of a stressor, which in this context is the jammer.

Antifragility is a concept popularized by Nassim Nicholas Taleb and is a term he coined in his 2012 book *Antifragile* [2]. Antifragility refers to systems that increase in capability, resilience, or robustness as a result of mistakes, faults, attacks, or failures. As Taleb explains in his book, antifragility is fundamentally different from the concepts of resiliency (the ability to recover from failure) and robustness (ability to resist failure). "The resilient resists shocks and stays the same; the antifragile gets better" [2].

In this paper, we apply the concept of antifragility to wireless communication systems. Specifically, we seek to exploit the presence of a jammer to achieve a communications gain relative to a jammer-free case. This should not be confused with anti-jamming, which seeks to mitigate jamming and perform at (or near) interference-free capability through the duration of an attack. Antifragile communications take mitigation one level further, by providing a boost during the attack. Additionally, our strategy should not be confused with self-jamming, i.e., friendly jamming, in which a secrecy channel is formed by *intentionally* transmitting noise along with the communication signal in order to prevent the eavesdropper from inferring any information.

The contributions of this work are summarized as follows:

1. Development of a novel *antifragile waveform*, used for manipulating a reactive jammer into relaying information, whereby a gain (relative to a non-jammed case) in throughput, connectivity, or covertness can be achieved.

2. Introduction of a generalized model for reactive jamming, applicable to both repeater-based and sensing-based jamming behaviors.

The remainder of this paper is organized as follows. Section II discusses background information on both antifragility and jamming. Section III introduces the antifragile scenarios that define the scope of this paper, while Section IV defines the jammer models under consideration. Section V introduces the components of an antifragile waveform, with emphasis on the jammer piggybacking strategy. Section VI develops theoretical channel capacities under each jammer model when using the antifragile waveform. Through numerical evaluation, Section VII provides feasibility regions for the proposed techniques (i.e., regions in which an antifragile gain occurs for the given scenario). Section VIII concludes.

3 Background

3.1 Related Work

With the popularization of the antifragility concept has come related literature in various fields. Antifragility is applied to the field of biology in [3], where proteins with flexible regions undergo functional alteration of their side residues or backbone. The authors of [4] propose antifragile "open systems", those that continuously communicate and interact with other systems outside of themselves. Safety engineering, and its connection to antifragility, is investigated in [5]. The authors of [6] present examples of antifragility being applied to engineering systems, although the authors do not describe any communications related examples. In the wireless domain, antifragility has been studied in complex networks for analyzing connectivity. In [7], the authors propose a new metric for complex networks used to quantify the impact of each node or failure in network connectivity. In terms of the open systems interconnection (OSI) model, the analysis in [7] occurs on the network layer while our analysis is primarily concerned with the physical and data link layers.

In [8], examples of antifragility in electronic systems are given, including the multipath phenomena in radio frequency (RF) transmission. The author points out how for the first half of the last century, multipath phenomena was thought of as harmful. Thus, multiple-input multiple-output (MIMO), which uses multipath to enhance system performance, can be considered an example of antifragile communications. Considering how comprehensive literature on MIMO is, our investigation into antifragile communications focuses on other aspects, such as jammer exploitation.

While not meant to be antifragile, research involving the establishment of a timing channel to counter a reactive jamming attack, including [9] and [10], has similarities with the approach described in this paper. The strategy in [9] is similar to our *replace with noise* jammer exploitation approach, in terms of using the presence or absence of a signal to carry information. Such a strategy does not attempt to evade the jammer, but rather, function *in spite* of jamming, although the authors do not go as far as to *exploit* the jammer. However, the work in [10] could be considered as jammer exploitation. In an approach similar to jammer piggybacking, the authors investigate the use of a timing channel to counter reactive jamming, where information is encoded in the interval between the instant when the jammer terminates its jamming signal and the beginning of the transmission of the next packet. The end of the jamming signal is used as a reference because the end of the previous packet is covered by the jammer's interference. While the approach in [10] involves the use of the jamming waveform to convey information, it is done using a different process and for a different reason than the approach described in this paper.

To the best of our knowledge, this is the first literature that specifically applies the concept of antifragility to wireless communications. In addition, this is the first literature that describes a method for piggybacking off a reactive jammer, to achieve a communications gain relative to a jammer-free case. To broaden the application of the techniques introduced in this paper, we have explored them under a wide variety of jamming behaviors. However, we make no claims as to the existence of these jammers in modern warfare.

3.2 Antifragility Compared to Similar Terms

To make the concept of antifragility more clear, we compare it with similar concepts.

Antifragile vs. Robust/Resilient: The concept of antifragile is most often confused with robustness or resilience. While an object or system can certainly be robust against stressors or harsh

Approved for public release; distribution is unlimited.

conditions, it is only antifragile if it benefits from them. In other words, compared to an established baseline performance of a given metric, a resilient system will never increase over the baseline due to harsh conditions, while an antifragile system has the ability to.

Antifragile vs. Adaptive: An adaptive system is one that changes its behavior based on information available at time of utilization (as oppose to having the behavior defined during system design). While adaptive systems allow for robustness under a variety of scenarios, they are not necessarily antifragile. In fact, it is often difficult to determine if an adaptive system is also antifragile, because baseline performance in an adaptive system is usually not well-defined, and typically based on the current conditions which may be changing.

Antifragile vs. Cognitive: We usually think of cognitive systems as incorporating artificial computational processes that act like a person. A cognitive entity is one that is capable of decision making, carrying out actions depending on its own goals and its perception of the world, and learning from experience. In the wireless domain, cognitive is most often discussed in the context of cognitive radios, which are able to perform well across a wide variety of harsh conditions. While *cognitive* characterizes how a system works, *antifragile* is more concerned with the output or performance of the system. An antifragile system can certainly contain cognitive components to it, but a cognitive system is only antifragile if it is able to satisfy the criterion of increasing in capability in some way *as a result of* a fault, attack, failure, or any negative condition.

3.3   Brief Jamming Taxonomy

Jamming threats can first be broken down into the type of system being targeted, the common three being: communications, radar, and radio navigation (e.g., Global Positioning System (GPS) and Global Navigation Satellite System (GLONASS)). Within the scope of wireless communication jamming, there are several different jamming strategies that could be used. Following the structure of the taxonomy proposed in [11], we classify jammers based on their key capabilities. Figure 1 shows the four major communications jamming capabilities and how they relate.
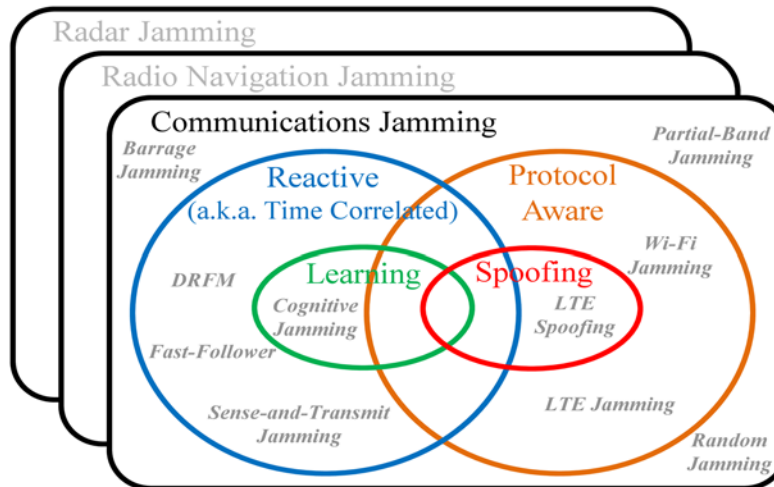

Figure 1: Venn diagram showing key capabilities of communications jammers, and how the capabilities relate in terms of feasible jammers. In each category, example jammers are shown in grey text.

Basic jammers, such as barrage or random jammers, operate as programmed and do not change their behavior. Thus, basic jammers do not have any of the four capabilities shown in Figure 1. Because basic jammers do not require a receiver, the communications system being jammed has no realtime influence on the jammer.

Reactive jamming, also known as time-correlated jamming and follower jamming, is a broad category of jammers that react to sensed energy or signals. Reactive jammers must have receiving capability, and a common form of reactive jamming is repeater jamming (a.k.a. digital RF memory or digital RF memory (DRFM) jamming) in which the jammer retransmits the signal it receives with a possible transformation applied [11]. Of the more complex types of jammers, reactive jammers are becoming more prevalent due to low-cost software-defined radios.

There are many reasons why an adversary would use a reactive jamming strategy. A communications link using frequency-hopping spread spectrum (FHSS) hops around in frequency very quickly, mitigating narrowband interference. If the jammer has no knowledge of the frequency hopping pattern, then the best it can do is transmit across the entire hopping band. This leads to the FHSS link gaining an anti-jam advantage equal to the FHSS processing gain. However, a jammer can overcome this anti-jam advantage by following the signal in frequency, which requires receiving the signal and retransmitting it as fast as possible. Outside of FHSS, reactive jamming can be used to surgically jam frames in a contention-based media access control (MAC) protocol, and avoid jamming an empty channel. This lowers the power consumption of the jammer, as well as makes the jammer harder to detect, because it is effectively hidden under an actual communications signal. If the reactive jammer wants to avoid retransmitting samples it receives, then it can use a sense-and-transmit approach, where it transmits internally generated noise immediately after sensing a signal. In the FHSS case this also requires detecting the current frequency channel the target is occupying. This detection process takes time, typically at least 10 to 100 symbols worth of observations [11]. Literature based around this sense-and-transmit reactive jamming strategy includes [12-14].

As depicted in Figure 1, a subset of reactive jamming is a jammer with the ability to learn. This type of jammer is one that can modify its behavior in real-time in response to its experiences, such as instances of successful or unsuccessful jamming actions [15]. On the other hand, protocol aware jamming is simply jamming against a particular wireless protocol, such as Wi-Fi or long-term evolution (LTE), typically by targeting a specific mechanism of the protocol instead of jamming the entire signal [11]. Spoofing is a subset of protocol aware jamming that requires transmitting a fake signal meant to masquerade as the real one.

While there are many categories of jamming, this paper focuses on the reactive type.

4   Methods, Assumptions, and Procedures

4.1   Motivating Scenario: Jammer Piggybacking

One approach to antifragile wireless communications is to manipulate a hostile jammer into unknowingly helping with the transmission of data. We call this strategy "jammer piggybacking". With jammer piggybacking, we seek to exploit a jammer that is correlated with the signal it is targeting, i.e., a reactive jammer. In Section 6 we show that information transmitted between two nodes can piggyback off the jammer, which effectively makes the jammer a relay node.

Figure 2 shows the geometrical configuration of a reactive jamming scenario. From an information theoretic point of view, if the transmitted signal *x* is jammed by signal *j*, an antifragile gain is only possible if the mutual information between the two is greater than zero, denoted as $I(x;j)>0$, where $I(a;b)$ indicates the mutual information between two random variables *a* and *b*. Mutual information helps to measure how much knowing one variable reduces uncertainty about the other. In other words, if the inequality is greater than zero, then the jamming signal contains information that the source node is sending to the destination node. This information could even be manifested in the presence or absence of the jamming signal, similar to how in On-Off Keying (OOK) information is conveyed by the presence or absence of a carrier. To perform jammer piggybacking, the destination node must be able to separate the two signals, demodulate them, and then combine them (or simply use the one that has the most integrity).

Figure 2: The geometrical configuration of a reactive jamming scenario illustrates the three channels involved. At each point a signal is received, a Signal-to-Noise Ratio (SNR), denoted with γ, is shown.

While under some scenarios it is possible to achieve a throughput gain from jammer piggybacking, with respect to a non-jammed case, an antifragile gain could also stem from other advantages. For example, if a low bit-rate signal known by the destination node is successfully relayed through the jammer, the destination node can use this signal to aid in null-steering (assuming it has that capability). By having a signal to correlate to, as opposed to just noise, the destination node can more quickly and accurately point a null towards the jammer. Another example involves solving the hidden node problem by manipulating a high-power jammer into providing a control channel (or rather a relay for a control channel) among a network that may include hidden nodes. A directional antenna should be used by the manipulating radio, so that the other radios only receive the jammer's version of the signal, and the original signal does not interfere with other transmissions. Both of these examples show how jammer piggybacking can still provide an antifragile advantage, even without achieving a higher-than-baseline bit-rate.

This concept of reactive jammer piggybacking is further developed in this paper.

4.2   Other Example Antifragile Strategies

In addition to jammer piggybacking, we propose the following methods of antifragile communications:

4.2.1   Achieving Coarse Time Synchronization

A high power pulsed enemy jammer (or even enemy radar) could act as a way for multiple radios to achieve coarse time synchronization. For example, if a group of radios in the same general area use time division multiple access (TDMA), slot transitions could be governed by the pulses of the

enemy transmitter. Therefore, this strategy allows a wireless network to increase in synchronization capability as a result of an attack.

4.2.2   Inducing Jammer Friendly Fire or Jammer Herding

Causing an enemy jammer to jam other systems friendly to it could have a clear antifragile advantage, since the total electronic attack performance increases as a result of the attack. Realizing this type of strategy would be based heavily on the specific communications protocol and jammer behaviors, and thus is beyond the scope of this paper and is left for future investigation.

4.2.3   Hiding in an Enemy Jammer's Signal

Low Probability of Intercept and Detection (LPI/LPD) is an important aspect of wireless communication waveforms meant for mission-critical use. This antifragile strategy involves increasing LPI/LPD by hiding the desired signal within a jamming signal. Clearly this is only possible if there is a way to cancel out the jamming signal at the destination node, leaving only the desired signal to be demodulated. In order to be truly hidden under a signal, the two signals must (mostly) overlap in time and frequency, and the *cover signal* must be, in general, higher power. One possible method of removing the jamming signal is to exploit cyclostationarity in the jamming signal, which could be a result of the jammer using a specific modulation scheme or using a repetitive pattern (such as a chirp signal). By hiding under the enemy jammer's signal while maintaining communications, LPI/LPD is increased as a result of an attack.

4.3   Defining Three Classes of Antifragility

Because antifragility is measured by an increase in some sort of performance, we will define three classes of antifragility, as shown in Figure 3. The y-axis is left unlabeled, as there are several different ways a communication system can increase in performance or capability (e.g., throughput, spectral efficiency, node connectivity). Class I antifragile systems increase in performance during an attack, but the antifragile gain does not persist once the attack ends. Most of the discussion in this paper involves Class I strategies. Class II systems, on the other hand, increase after the attack (usually due to information obtained during the attack). Note that in Figure 3, the Class II example shows what is essentially perfect resiliency during the attack, which may not always be the case (throughput may even drop to zero during the attack). Class III systems represent a combination of both Class I and Class II, in which the antifragile gain occurs at the beginning of the attack and persists after the attack ends.

Figure 3: Diagram of antifragile, resilient, and fragile systems. Antifragile systems are able to maintain objective performance through a period of adverse conditions and achieve super-objective performance during or after the period of adverse conditions.

5 System Model

In this section we describe the channel model and jammer models under consideration. As mentioned before, we are most interested in jamming of a reactive nature, with the goal of making the jammer act as an unwitting relay. Throughout this paper we refer to the transmitter as the source node, and the receiver as the destination node.

5.1 Channel Model

We consider a transmitted signal x that goes through a channel h with propagation delay τ_h. This signal x can be eavesdropped on by the jammer via a channel k_e with delay τ_e, and then jammed by signal j that goes through a channel k_j with delay τ_j. We assume all three channels are noisy memoryless channels, thus h, k_e, and k_j correspond to channel coefficients, which may or may not be time-varying. In addition, the parameter τ_{jam} denotes the delay due to the jammer's RF chain and any sensing that may be performed. Noise n_d and n_j is the additive white Gaussian noise seen at the destination node's receiver and jammer's receiver respectively, with variance σ_d^2 and σ_j^2. These signals and channels are depicted in the generalized reactive jamming model, shown in Figure 4. Under this generalized model, $\beta(t)$ represents a time-varying transform applied to the received signal at the jammer, while $w(t)$ represents an internally generated jamming waveform.

When the jammer uses a behavior that only involves repeating the received signal with a transform $\beta(t)$ applied, $w(t)=0$ and the received signal y is given by:

$$y(t)=hx(t-\tau_h)+k_j\beta(t)\left[k_e x(t-\tau_e-\tau_{jam}-\tau_j)+n_j\right]+n_d \tag{1}$$

Likewise, when the jammer uses a sensing-based behavior that involves transmitting a jamming waveform $w(t)$, $\beta(t)=0$ and the received signal y becomes:

$$y(t)=hx(t-\tau_h)+k_jw(t-\tau_e-\tau_{jam}-\tau_j)+n_d \tag{2}$$

Note that while the sensing-based jammer does not retransmit x or n_j, the eavesdrop delay τ_e is still a factor because the jammer has to perform sensing.

The conditional probability of y given x and j, denoted as $p_{Y|X,J}(y|x,j)$ is assumed to be stationary and a function of the communications channels, which we will not impose a model for. The marginal distribution $p_X(x)$ is determined by the transmitter's physical layer parameters such as modulation and coding scheme.

Because the reactive jammer acts as a form of memory, we can apply Shannon's formula for capacity through a channel with memory, which is stated as [16]:

$$C=\lim_{n\to\infty}\sup_x\frac{1}{n}I(x^n;y^n) \tag{3}$$

Let C_1 and C_3 be the channel capacity before and after the jamming attack respectively, as labeled in Figure 3. In both of these cases $j(t)=0$ and the traditional memoryless channel capacity applies. Let C_2 be the capacity during the jamming attack, when $j(t)$ takes on a certain behavior as discussed in the next subsection. Therefore, the criterion for Class I antifragility in the jammer piggybacking context is simply $C_2>C_1$, indicating an increase in capacity as a result of jamming. Likewise, the criterion for Class II antifragility is $C_3>C_1$. Class III antifragility requires that both inequalities be true, as depicted in Figure 3.

Figure 4: Generalized model for reactive jamming.

In Section 7 these capacities will be defined in terms of the SNRs of the three different channels involved in the analysis. For ease of analysis, we assume that all three SNRs are constant. The SNR of the source node signal as seen by the destination node is denoted as γ_s. The SNR of the source node signal as seen by the jammer when eavesdropping is denoted as γ_e. Lastly, the SNR of the jamming signal (not taking into account noise n_j) as seen by the destination is denoted as γ_j. These three SNRs are defined as follows:

$$\gamma_s=\frac{E\left[|x|^2|h|^2\right]}{\sigma_d^2} \qquad \gamma_e=\frac{E\left[|x|^2|k_e|^2\right]}{\sigma_j^2} \qquad \gamma_j=\frac{E\left[|j|^2|k_j|^2\right]}{\sigma_d^2} \tag{4}$$

All parameters associated with the channel and jamming model are summarized in Table 1.

Table 1: Summary of model parameters.

Symbol	Description
h, τ_h	main channel coefficient and delay
k_e, τ_e	eavesdrop channel coefficient and delay
k_j, τ_j	jamming channel coefficient and delay
τ_{jam}	jammer's RF + sensing + intentional delay
n_j	channel noise at the jammer's receiver
n_d	channel noise at the destination node's receiver
$\beta(t)$	time-varying transform applied to signal
$w(t)$	internally generated jamming waveform
$\gamma_s, \gamma_e, \gamma_j$	SNR of main, eavesdrop, and jamming channel respectively

5.2   Reactive Jamming Models and Behaviors

We will model the transmitted signal $x(t)$ as a single-carrier signal with a certain carrier frequency f, phase φ, and amplitude A:

$$x(t)=A cos(2\pi ft+\varphi) \tag{5}$$

Using this approach, these jammer models can apply to a communications link that uses FHSS, and also a traditional single-carrier based signal. We also assume the jammer does not know values of x a priori.

The basic behavior of a reactive/repeater jammer is that it receives a signal transmitted from a target transmitter and retransmits it with a possible transform applied, in a manner intended to jam the target receiver. By repeating the target signal, the jammer can follow the two radios as they hop around in frequency, countering the protection associated with FHSS (which could be over 20 dB). In order for the jammer to increase its effectiveness, it tries to transmit the jamming signal with no significant frequency offset, in order to better align with the target signal in frequency. Our first jammer model under consideration, the DRFM, is one that simply retransmits the target signal on a sample-by-sample basis [17, 18]:

DRFM: The jammer retransmits the received signal with a constant amplification gain β_A, such that $\beta(t)=\beta_A$:

$$j(t)=\beta_A \left[k_e x(t-\tau_e-\tau_{jam})+n_j \right] \tag{6}$$

This jamming method does not require any form of frequency detector, but the jammer's bandwidth must be sufficiently high to cover the hopping range of the target link (if FHSS is in use). In the case that the noise power is not significantly higher than the repeated signal power, this jamming model can be thought of as an amplify-and-forward relay.

In some cases the jammer may want to transmit internally generated noise in place of the target signal. We generalize this jamming behavior with the following model:

Replace with noise: The jammer transmits internally generated random noise on the frequency that x was received on. When targeting a FHSS signal, this type of jammer is referred to as a follower jammer [18]. In terms of the generalized reactive jamming model, $w(t)=n_{jam}(t)$ when signal x is detected, and the jamming signal is given by:

$$j(t)=n_{jam}(t-\tau_e-\tau_{jam}) \tag{7}$$

when x is detected, and $j(t)=0$ otherwise. In most cases $n_{jam}(t)$ can be modeled as a zero-mean band-limited Gaussian random process with variance σ_{jam}^2, spanning the bandwidth of signal x. This method requires some form of an energy detector to detect which frequency the target is transmitting on and if the signal is present. The delay associated with this detection process is included in τ_{jam}. Literature using this jammer model in the wireless communications domain includes [12-14, 18, 19].

While the *replace with noise* model seems like an effective method for jamming, in some scenarios the jammer may not be able to detect the frequency quick enough to successfully jam the enemy. Therefore, we will investigate some alternative behaviors that the jammer could use to prevent being an amplify-and-forward relay, without having to perform signal detection.

Phase flipping: The jammer randomly flips the phase of the received symbols:

$$j(t)=U(t)\left[k_e x(t-\tau_e-\tau_{jam})+n_j\right] \tag{8}$$

where $U(t)$ is a pseudorandom sequence drawn from the set $U(t)\in\{1,-1\}$ with probability mass function given by $f_U(u)=0.5:U=1,-1$. In terms of the generalized reactive jamming model, $\beta(t)=U(t)$. This could be realized by multiplying the received signal by a square wave that alternates between +1 and -1 randomly. This model is especially useful because practical modulation schemes, such as Phase-Shift Keying (PSK), Quadrature Amplitude Modulation (QAM), and Amplitude-Shift Keying (ASK), can be corrupted through this approach.

Realistically, a jammer would have no reason to distinguish between symbols (symbol-level timing synchronization is not trivial), when it can simply change the value of $U(t)$ often enough to corrupt data carried in the phase. However, we use this model as a way to generalize the jamming tactic. The drawback to this jamming technique is that it spreads the transmitted signal in the frequency domain, which means there will be energy that does not overlap with the target signal. To reduce spreading, the jammer must increase the time between flipping the phase. A favorable switching period for the jammer would be one that is more frequent than the presence of equalization pilots, so that the phase shifts do not get equalized out at the destination node, but not so quick that the signal gets overly spread in frequency.

An alternative to the phase flipping approach would be to modulate the amplitude.

Modulate amplitude: Identical to the previous tactic, except the jammer randomly modulates the amplitude, and $\beta(t)=V(t)$:

$$j(t)=V(t)\left[k_e x(t-\tau_e-\tau_{jam})+n_j\right] \tag{9}$$

where $V(t)$ is a pseudorandom sequence of positive numbers. An example probability density function for V, in which the signal's average power is conserved, is $f_V(v)=0.5$ where $0\leq v\leq 2$.

If a jammer modulates the amplitude between a negative and positive value, we will consider it as also modulating the phase. In this *phase and amplitude modulating* case, the jamming signal likely resembles noise, and the *replace with noise* model is most appropriate.

Lastly, we will consider a model similar to *replace with noise*, except the jammer replaces the signal with a continuous wave (CW).

Replace with CW: The jammer transmits a CW on the frequency that the target signal was received on (i.e. $\beta(t)=0$ and $w(t)$ is a sinusoid). Like the *replace with noise* model, this model requires some form of an energy detector. The phase and amplitude of the transmitted CW is independent of the target signal, leaving φ and A out of the equation for $j(t)$:

$$j(t)=cos(2\pi f(t-\tau_e-\tau_{jam})) \tag{10}$$

when x is detected, and $j(t)=0$ otherwise.

Table 2 summarizes the discussed jamming models in the context of the generalized reactive jammer model shown in Figure 4.

Table 2: Jammer models in the context of the generalized reactive jamming model.

Jammer Model	$\beta(t)$	$w(t)$
DRFM	β_A	0
Replace with Noise	0	$\mathcal{N}(0,\sigma_{jam}^2)$
Phase Flipping	$U(t):u\in\{1,-1\}$	0
Modulate Amplitude	$V(t):v\geq0$	0
Replace with CW	0	$cos(2\pi ft)$

Nulling-type attacks are not considered in this paper, because the accuracy of channel state information required by the jammer makes them infeasible [11]. Additionally, a time-delay type transformation is not included as a jammer model because it would be equivalent to the DRFM model with τ_{jam} intentionally increased. While the presented jammer models are not all-exhaustive, we feel they represent a significant portion of possible reactive jamming strategies.

6 Components of an Antifragile Waveform

In this section we discuss components of an antifragile waveform specific to the reactive jammer piggybacking strategy, although many of the components discussed also apply to the other antifragile strategies in Section 4.2. The intent of an antifragile waveform is not to provide a radio with an antifragile gain at all times. Rather, an antifragile waveform grants the ability to achieve an antifragile gain when the situation allows for one. A functioning antifragile strategy is based on the scenario at hand, such as the type of communications link, the reactive jammer's characteristics, and the delays involved. An effective strategy involves implementation of a series of different schemes that exploit the jammer, and a classifier that can identify the scenario at hand and estimate parameters associated with it. Lastly, it must incorporate an engine that can assign the most effective antifragile (or mitigation) scheme to the given scenario. Figure 5 illustrates the

components of the proposed antifragile system (highlighted in yellow) and how they fit into a communications system; each subsection in this section corresponds to one of the yellow boxes.

Figure 5: The components of the proposed antifragile system (highlighted in yellow) and how they fit into a communications system.

6.1 Delay Estimator

The most important question to answer is, "what is the delay associated with the reactive jammer?" In other words, the delay between when the actual signal is received and jamming signal is received at the destination node. We will denote this delay as d, and it is given by $d=\tau_e+\tau_{jam}+\tau_j-\tau_h$. We will split the possible delay scenarios into three cases, where T_{sym} is the period of one symbol, and T_{hop} is the period of one hop (or one packet/frame). These three cases are listed below and depicted in Figure 6.

1. $d<T_{sym}$

2. $T_{sym}\leq d<T_{hop}$

3. $d\geq T_{hop}$

Figure 6: Examples of how the source and jamming signal are received at the destination node, for each of the three time delay cases.

The first case in which $d{<}T_{sym}$ represents a jammer that is close enough and has a low enough delay to overlap on a symbol-by-symbol basis. The third case in which $d{\geq}T_{hop}$ represents a slow repeater jammer that does not cause any link degradation, but can still be exploited.

To initially estimate delay d (before the jamming behavior has been classified), the destination node can simply observe the time lag between the hop/frame preamble, and the next burst of energy received. While it may take a few hops to get an accurate initial estimate, this process can be performed simultaneously with any other jammer detection on the platform. We will denote the current estimate of d as $\hat{d}$. Once the jammer is actively relaying information, updates to $\hat{d}$ can be made at the destination by observing the delay between two copies of any given preamble.

While the delay is an important measurement, it does not capture the time it takes the jammer to stop transmitting a jamming signal after the received signal at the jammer ends. We will denote this value as the *ending-lag*, and it is depicted in Figure 7. In cases such as an analog repeater jammer, this value of *ending-lag* would be near-zero because the jammer is simply retransmitting what it receives. However, a digital repeater jammer that uses sense-transmit cycles could have a short delay before it stops transmitting.

Figure 7: Depiction of *ending-lag*, which we define as the delay the jammer takes to stop transmitting after it stops receiving a signal.

6.2 Jammer Classification

Next, the radios must classify the reactive jammer's behavior. We assume that there has been time-frequency orthogonality established between the desired signal and jamming signal, a process we discuss in the next subsection. Under this assumption, the destination node can separate the two signals, extract features from the jamming signal, and perform jammer classification. The classification results must be relayed back to the source node. Therefore, it is assumed that there is, at a minimum, a low data rate channel that the destination node can use to share the classification results with the source node (typically referred to as the feedback or return channel). The number of bits associated with classification results can be reduced using a predefined lookup table. A probing function is used to send known symbols (a.k.a. pilots) that are designed to help distinguish among the various jamming models discussed in Section 5. We propose using machine learning classification, where the features used for classification consist of:

Demodulation of Pilot Symbols: a straightforward method of detecting if the strategy *phase flipping* or *modulate amplitude* is in use. Using QAM modulation on a selection of the pilot symbols, such as 16-QAM, would allow the destination node to determine if the jammer is

modulating the phase (the amplitude portion of the constellation would remain intact) or modulating the amplitude (the phase information would remain intact).

Crest Factor (CF): a feature of the jamming signal in the frequency domain, which can be used to differentiate between *replace with CW* and the other models. CF indicates the ratio of peak value to the quadratic mean: $CF=|x|_{peak}/x_{rms}$. Note that peak-to-average power ratio (PAPR), an important parameter in wireless communications, is the square of the CF.

Cross-correlation: simply the correlation of *y* with *y* after rough timing realignment, in which a strong peak would indicate the *DRFM* model is in use. A cross-correlation (i.e., sliding dot product) is used instead of just one dot product to take into account the fact that $\hat{d}$ is likely only accurate enough for rough timing realignment, not down to the symbol or sample level.

While none of the above features explicitly detect the *replace with noise* jamming model, we can treat this model as the *default* jamming behavior that is assumed if none of the other jamming models seem to be appropriate for the scenario. This makes sense because the *replace with noise* model makes the fewest assumptions about the jammer's behavior, thus following Occam's razor.

Based on extracting these three features, the classifier can determine the specific model that most closely matches the jammer's behavior. Multiple hops worth of observations should be used for accurate classification, and features must be extracted from each hop individually. The transitory phase due to delay estimation and jammer classification may span several hops, but it is likely insignificant compared to the period of time the reactive jammer is active. The specific classifier we suggest using is the Support Vector Machine (SVM), due to its superior performance across a large range of applications [20] (for a comparison between classification methods, we refer the reader to [21]). Training is performed offline, using several instances of each jammer model (either simulated or implemented in hardware). Retraining could also be performed in the field if a new threat is discovered, or to take into account minor variations to the existing threats.

6.3   Forcing Orthogonality

To avoid the transmitted signal and jamming signal being received co-channel (overlapping in frequency and time), we must force orthogonality using observations of the jammer's delay and *ending-lag*. The process of causing orthogonality between a desired signal and a jamming signal is at the heart of anti-jamming, and as such, we make use of common anti-jamming strategies as part of the antifragile waveform. We assume that the communications link under study is able to null certain symbols (e.g., by assigning the value $0+0j$ to them). The proposed solution to forcing orthogonality is based on which of the three delay cases occur.

When $d<T_{sym}$, true time-frequency orthogonality is simply not possible because symbols are overlapping, causing co-channel signals. However, if the Jammer-to-Signal Ratio (JSR) is significantly high, then orthogonality is not needed to achieve an antifragile gain, as combining is not necessary and the destination node can use the jamming signal and treat the actual signal as noise. Likewise, if a phase array antenna can isolate the two signals, then orthogonality is achieved *spatially*.

When $T_{sym}{\leq}d<T_{hop}$ and the radio is hopping as fast as it can, it must null its *e* last symbols of each hop or frame, where:

$$e = \frac{T_{hop}-d}{T_{sym}} \qquad . \tag{11}$$

15

In other words, the source node refrains from transmitting on symbols that would overlap with the jamming signal. If the fraction of overlapping signals is large, methods from the previous delay case can be used.

When $d \geq T_{hop}$, orthogonality can be created by hopping in such a way that there are no *collisions* in time and frequency on a hop-basis. With this approach, the actual signal and jamming signal are received orthogonally, and the signals can be separated and demodulated independently.

6.4 Modulation Scheme for the Antifragile Waveform

The fundamental thrust of any jammer-based antifragile communications method is to exploit the signalling dimensions enabled by the jamming system's transmissions, which may manifest in time, frequency, or space. Confining ourselves in this discussion to scenarios where the transmitted signal and jamming signal are received orthogonally, the optimal modulation scheme is intimately tied to which dimensions the jammer transform function either preserves from the transmitted signal or introduces to potential signalling by its operation. Consequently, an accurate jammer model (see Section 5) is critical to realizing a useful communications link, and effective classification methods are necessary to exploit encountered systems.

Several common jammer types do not preserve incoming signal amplitude and phase information, instead transmitting a newly generated signal such as a noise-like waveform or CW tone in the channel. If the jamming signal is sufficiently narrowband and its duration consistently deterministic, the source transmitter can effect a noncoherent pseudo-frequency shift keying (FSK) relay link by modulating the frequency of its own transmissions. Similarly, a frequency-preserving jammer transform may support chirp modulation types. Alternately, a consistently deterministic jammer duration alone enables noncoherent modulation schemes such as straightforward OOK or the family of Pulse Position Modulation (PPM) types, which may accommodate differential coding, overlapping, or other features. An example of PPM transmitted *through* the jammer is shown in Figure 8, with the *ending-lag* highlighted in red.

Figure 8: Example waveform when using PPM. Each tick represents a symbol period, with the red portion of each symbol representing the *ending-lag*. The actual signal and jamming signal have been aligned in absolute time (i.e., *d*=0) for the sake of presentation.

Depending on the jammer architecture, the system may channelize its wideband received input in order to monitor multiple subchannels at the same time. Therefore, if the antifragile radio can determine the specific subchannel configuration used by the jammer, it has the potential to achieve simultaneous modulation on multiple subchannels, resulting in an overall transmit scheme that resembles Frequency Division Multiplexing (FDM). The radio must simply provide a guard band between exploited subchannels.

The phase flipping jammer model discussed earlier may disrupt conventional phase signalling schemes. One solution is to avoid phase information entirely by using what we call "Positive-ASK", which is ASK with constellation points only in the positive half of the real axis. Postive-

ASK can be thought of as the ASK equivalent of single-polarity Pulse-Amplitude Modulation (PAM). Conversely, for the relatively unusual amplitude-only modulating jammer, a phase modulation approach (e.g., 8PSK) would work.

Because it represents an amplify-and-forward relay, the DRFM jammer model is the simplest to deal with, requiring no change in the original waveform.

A summary of the modulation scheme assignments for each jammer model is given in Table 3.

Table 3: Modulation assignments for each jammer model.

Jammer Model	Corresponding $j(t)$	Modulation Scheme
DRFM	$\beta_A(k_e x + n_j)$	Default Scheme
Replace with Noise	$n_{jam}(t)$	FSK, chirp, OOK, PPM
Phase Flipping	$U(t)\beta(k_e x + n_j)$	Positive-ASK
Modulate Amplitude	$V(t)\beta(k_e x + n_j)$	PSK
Replace with CW	$cos(2\pi ft)$	FSK, chirp, OOK, PPM

6.5   Combining Technique

Since the goal of this antifragile scheme is to relay some or all of the source node's data using the jammer, the destination is presented with two different signals carrying the same information. Thus, we have a similar situation to receive diversity in MIMO communications, with the exception that the two signals are likely received at different power levels. Within diversity combining there are three common techniques: selection combining (SC), maximal-ratio combining (MRC), and equal gain combining. In SC, the receiver simply selects the signal that is received with the highest SNR (or other channel quality indicator). MRC weights each received signal before combining, in such a way that the combined SNR is maximized (thus being the optimal scheme in terms of SNR). It has been shown that when using MRC, the combined SNR is simply the summation of the individual SNRs when in linear form [22]. Equal gain combining is a special case of MRC in which the weights applied to each signal are equal (usually set to 1). Analysis of these combining techniques with unequal SNR is performed in [22]. Because the SNR of the actual signal and the effective SNR of the signal relayed through the jammer could vary greatly, equal gain combining does not make any sense for this application. MRC and SC are compared in Figure 9, under a constant main channel SNR of 5 dB and a varying JSR. To achieve an antifragile gain in this example, the SNR after combining must be above 5 dB. It is assumed that the channel noise at the jammer's receiver is the same as the noise at the destination node, and fading is not taken into account.

Also included in Figure 9 is a case where orthogonality is not achieved, so instead of combining the two signals, the destination node simply decodes the jamming signal and treats the actual signal as noise. In this case, the JSR must be at least 7.5 dB for the strategy to function.

Figure 9: Average SNR after combining, for maximal-ratio and selection combining, when the main channel SNR is a constant 5 dB. Also included is the SNR when only decoding the jamming signal, treating the actual signal as noise, which would only make sense if orthogonality is not achievable and JSR is high.

7 Theoretical Channel Capacities

We will now find the theoretical channel capacity under each jammer model described in Section 5.2, when using the jammer piggybacking antifragile strategy. The channel capacities in this section are primarily provided to demonstrate the different parameters associated with the process of piggybacking off each jammer type, and how throughput performance is influenced.

The channel capacity when no jammer is present, which will act as a baseline to determine when the antifragile requirement is met, is given by Shannon's theorem [23]:

$$C = B \log_2 \left(1 + \gamma_s\right) \tag{12}$$

where C is channel capacity in bits per second and B is channel bandwidth in Hz (i.e., the passband bandwidth of the signal). In Equation 12, the signal is assumed to be a wide-sense stationary zero mean complex Gaussian random process.

For each jammer model, the channel capacities incorporate a factor that represents that fraction of orthogonality achieved, on a per-hop or per-frame basis (see Section 6.3 for more information on timing). We will denote this fraction as D_{orth}, where $D_{orth}=1$ corresponds to zero overlap in the actual signal and jamming signal, and $D_{orth}=0$ corresponds to full overlap. As such, this parameter appears as a multiplier in the channel capacity equations. Lastly, we assume MRC is the form of combining used.

The simplest jammer model to analyze is the **DRFM**, which resembles an amplify-and-forward relay. The effective relay channel (i.e., the channel through the jammer) is a combination of channels k_e and k_j. The SNR of this total link through the jammer, which we will denote as γ_{ej}, is based on the value of β_A. In relaying literature, it is assumed that the relay's amplification gain is adjusted in realtime to satisfy the relay's (or jammer's) maximum output power constraint, denoted as $|j|^2$ [24]. Under this condition, β_A will always equal:

$$\beta_A = \sqrt{\frac{|j|^2}{|x|^2 |k_e|^2}} \tag{13}$$

While this may not always be the case for a uncooperative jammer, it is not an unreasonable assumption, given that the jammer may also want to maximize its output power. Equation 13 allows us to approximate the SNR of the link through the jammer [25]:

$$\gamma_{ej} = \frac{\gamma_e \gamma_j}{\gamma_e + \gamma_j + 1} \tag{14}$$

When MRC is used, the overall SNR is simply $\gamma_s + \gamma_{ej}$. Therefore, the overall channel capacity for the DRFM case is:

$$C = D_{orth} B \log_2 \left(1 + \gamma_s + \frac{\gamma_e \gamma_j}{\gamma_e + \gamma_j + 1} \right) \tag{15}$$

When a **phase flipping** jammer is present, we transmit in a manner that does not involve the imaginary portion of the constellation, as discussed in Section 6.4. Therefore, we must use the channel capacity formula for a single dimension, instead of complex [26]. When using MRC, the resulting channel capacity becomes:

$$C = \frac{D_{orth} B}{2} \log_2 \left(1 + \gamma_s + \frac{\gamma_e \gamma_j}{\gamma_e + \gamma_j + 1} \right) \tag{16}$$

$$P_e = \frac{1}{k} \left[\sum_{i=1}^{M/2} \frac{w_i'}{2\pi} \left(\int_0^{\pi - \frac{2i\pi - \pi}{M}} \exp\left(-\gamma \frac{\sin^2[(2i-1)\pi/M]}{\sin^2 \Theta} \right) d\Theta - \int_0^{\pi - \frac{2i\pi + \pi}{M}} \exp\left(-\gamma \frac{\sin^2[(2i+1)\pi/M]}{\sin^2 \Theta} \right) d\Theta \right) \right] \tag{17}$$

Exploiting an **amplitude modulating** jammer involves using M-PSK, as discussed in Section 6.4. The bits per second when using M-PSK is simply $C = \log_2 M$. The maximum reliable value of M is a function of SNR, which can be approximated using the theoretical bit error rate for uncoded M-PSK [27], given in Equation 17. Finding an upper limit on C requires taking into account all combinations of M and channel coding schemes, which is not feasible. In Section 8, when feasibility regions are evaluated, we simply use Equation 17 and require the bit error rate to be above 10%.

The **replace with noise** and **replace with CW** jammers involve signal detection, and transmit an internally generated signal, $w(t)$, as the jamming waveform. Thus, the amount of information relayed through the jammer is partially based on the accuracy and speed of the jammer's detection process. If we assume the jammer has no a priori knowledge of the target waveform, then the only difference between observing a signal and observing noise is the statistical average energy they contain. Therefore, the optimum detector compares the average energy in an observed waveform to a threshold, also known as an energy detector or radiometer [28]. Probability of detection, P_D, of a Neyman-Pearson type energy detector is parameterized by eavesdrop SNR, γ_e, and number of samples, and given by [28]:

$$P_D = 1 - \Gamma\left(\frac{n}{2}; \Gamma^{-1}\left(\frac{n}{2}; 1 - P_{FA} \right) (1 + \gamma_e)^{-1} \right) \tag{18}$$

where $\Gamma(x,y)$ is the incomplete gamma function, P_{FA} is the probability of false-alarm, and n is the number of samples taken from the observed waveform.

The amount of information relayed through the jammer is also based on the symbol rate, which must be decided on by the source node, and is a function of the jammer's *ending-lag* (so that symbols through the jammer do not overlap, see Section 6.1 for more information). If we assume the symbol period is equal to the *ending-lag* plus a safety margin factor, which we will denote as T_{guard}, the resulting symbol rate is $(T_{end-lag}+T_{guard})^{-1}$. This safety factor can be used to account for any jitter associated with the jammer's delay d. In a practical system, there would be digital signal processing required to verify the symbol rate is usable.

We formulate the bit rate based on using OOK with N number of multiplexed signals, where N is an integer greater than zero. We must assume that the jammer's detection process is very accurate, else we cannot formulate an equation for the rate at which information can be *reliably* transmitted. If this is the case, the capacity is:

$$C=N \left(T_{end-lag}+T_{guard}\right)^{-1} \quad if P_D \approx 1 \quad and \quad P_{FA} \approx 0 \tag{19}$$

Simply put, the data rate is equal to the number of multiplexed streams multiplied by the symbol rate, but only if the jammer is able to accurately perform signal detection. While no SNR appears in Equation 19, it should be noted that γ_e is a parameter in Equation 18, while the other two SNRs must simply be high enough for reliable transmission of OOK, which is around 0 dB [29]. Equation 19 is valid for both the *replace with noise* and *replace with CW* jammer models, because when using OOK in such a manner, the bandwidth of the jammer's signal is not a factor (as long as the signal can be demodulated at the destination node). Rather, the symbol rate is a function of the jammer's timing.

8   Results and Discussion

In this section we provide numerical results showing the feasibility regions in which jammer piggybacking provides an antifragile gain. Feasibility is evaluated by comparing throughput with the baseline (jammer-free) case.

8.1   Simulation Scenario and Conditions

The simulation scenario matches Figure 2, with the jammer taking on each of the models described in Section 5 (although *replace with noise* and *replace with CW* are combined, because they provide the same results).

We vary JSR and main channel SNR, γ_s, instead of all three SNRs, for the sake of two-dimensional results. The only requirement is that we assume $n_j=n_d$, so that $JSR=\gamma_j/\gamma_s$ for a sensing jammer, and $JSR=\gamma_{ej}/\gamma_s$ for a repeating jammer. We have decided to use JSR as the x-axis and produce plots for a SNR of 3 dB and 10 dB.

Adaptive modulation and coding (AMC) is used, and we assume an additive white gaussian noise (AWGN) channel. For a given set of JSR and SNR, the best modulation scheme, modulation order, and code rate is evaluated (based on a pre-populated table). The only constraint is that the bit error rate after decoding must be below 10^{-3}. To provide numerical results that reflect a high-

Approved for public release; distribution is unlimited.

performance communication system, low-density parity-check (LDPC) codes with code rates spanning 1/6 to 16/17 are used. In terms of modulation schemes available to the AMC engine, *phase flipping* uses Positive-ASK, *amplitude modulating* uses PSK, and *replace with noise* uses OOK. Both the baseline (no jammer present) and *DRFM* cases can choose from PSK, QAM, and ASK.

8.2   Simulation Results

Figures 10 and 11 show simulation results under full orthogonality, which can occur when the delay is greater than one hop, or delay is less than one hop and there is spatial orthogonality. Each point at which the modulation and coding scheme changes is identified by sharp transitions in the curve. We remind the reader that the baseline in each plot refers to the communications link when there is no jammer present; the baseline modulation and coding scheme is purely based on the main channel's SNR, and thus is constant in each plot. The theoretical channel capacities derived in Section 7 are shown for the DRFM and phase flipping case, representing the upper bound on throughput given the waveform constraints (the other two capacities cannot be represented without additional assumptions).


Figure 10: Feasibility region of reactive jammer piggybacking when the main channel SNR = 3 dB. Each sharp increase corresponds to the modulation and coding scheme being adapted for a higher quality channel. A curve above the baseline indicates an antifragile gain, as highlighted in yellow for the DRFM case.

Figure 11: Feasibility region when the SNR = 10 dB.

For both levels of SNR, the DRFM case has the best performance, which is expected considering the lack of constraints on the waveform under this case. A loss is associated with having to use PSK (in response to an *amplitude modulating* jammer) and Positive-ASK (in response to a *phase flipping* jammer), although the loss is not large enough to prevent an antifragile gain altogether when SNR = 3 dB. Being forced to use OOK (in response to a *replace with noise* jammer) leads to a lack of an antifragile gain across the entire region, for all levels of SNR. However, these results reflect only having *one* data stream of OOK. As discussed before, it is likely possible (depending on the scenario) to transmit a large number of data streams in parallel using FDM.

It is expected that the antifragile gain mostly occurs in the right-hand portion of each plot, where the JSR is higher than zero. The antifragile gain when JSR is negative is largely due to the combining gain, while the gain when JSR is positive is due to the fact that the jamming signal reaches the destination at a much higher effective SNR than the main channel. Fortunately for this strategy, jammers typically intend to operate in the right hand portion of the plot, where JSR is above 0 dB (communications with adaptive modulation and coding tend to become denied at a SNR lower than 0 dB).

When full orthogonality cannot occur, such as when the delay d is between one symbol and one hop, a loss must be included to take into account the nulled symbols at the end of each hop. This loss simply shifts each curve down by a certain percent. Figure 12 shows an example scenario when SNR = 3 dB, and the actual signal and jamming signal overlap during 50% of each hop. The result is a 50% loss in throughput for each antifragile curve. It can be seen that the start of the antifragile gain shifts to the right, requiring higher jammer power compared to a fully orthogonal scenario.

In cases where an antifragile gain in the context of throughput is not feasible, it may still be possible to achieve antifragility by using the jammer to provide a low data rate control channel, as discussed at the end of Section 4.1.

Figure 12: Feasibility region when the SNR = 3 dB and there is only 50% orthogonality.

9 Conclusion

In this paper, we have introduced the concept of antifragile wireless communications, through a novel strategy that exploits a reactive communications jammer. Several reactive jamming models were described, including both repeater-based and sensing-based. We have outlined guidelines for realizing an *antifragile waveform*, and shown that an antifragile gain is possible under a wide variety of reactive-jamming scenarios.

As the sophistication of communications systems and jammers increases, reactive jamming will likely become a bigger threat in military and other mission-critical domains. Therefore, incorporating antifragility will not only improve protection of radios, but also bring about performance improvements in harsh conditions. In addition, a system that is able to achieve any level of antifragility disincentivizes an enemy jamming mission in its entirety, assuming that mission has nonzero cost. Even if the enemy is aware of its target's antifragile capability and prevents it from achieving an antifragile gain (e.g., by using basic barrage jamming), that antifragile capability indirectly leads to a less effective attack.

As part of future research we will continue to develop the concept of antifragile communications, including an investigation into how fading affects the performance of jammer piggybacking. There are several areas in which the concept can be expanded into future topics, including network-layer strategies, spatial signalling dimensions, reduction or elimination of required feedback signals, jammer herding, multi-source signalling, and interference alignment. In the case of the *replace with noise* jammer, the optimal symbol rate and number of frequency multiplexed signals is still an open question. Lastly, a deeper investigation into the antifragile schemes proposed in Section 4.2 could lead to additional applications.

References

[1] M. Lichtman, M. T. Vondal, T. C. Clancy, and J. H. Reed, "Antifragile communications," *IEEE Systems Journal*, vol. PP, issue 99, pp. 1-12, Feb 2016.

[2] N. N. Taleb, *Antifragile: Things that gain from disorder*. Random House LLC, 2012.

[3] A. Danchin, P. M. Binder, and S. Noria, "Antifragility and tinkering in biology (and in business) flexibility provides an efficient epigenetic way to manage risk," *Genes*, vol. 2, no. 4, pp. 998–1016, 2011.

[4] V. De Florio, "Antifragility = Elasticity + Resilience + Machine Learning Models and Algorithms for Open System Fidelity," *Procedia Computer Science*, vol. 32, pp. 834–841, 2014.

[5] E. Verhulsta, "Applying systems and safety engineering principles for antifragility," *Procedia Computer Science*, vol. 32, pp. 842–849, 2014.

[6] K. H. Jones, "Engineering Antifragile Systems: A Change In Design Philosophy," *Procedia Computer Science*, vol. 32, pp. 870–875, 2014.

[7] R. Melo, A. Santos, M. Nogueira, and D. Medhi, "Resilience and Knowledge in a Metric for Heterogeneous Wireless Connectivity," Technical Report RT-DINF 003/2013, University of Missouri-Kansas City, MO.

[8] R. W. Lucky, "Antifragile Systems [Reflections]," *IEEE Spectrum*, vol. 50, no. 3, pp. 28–28, 2013.

[9] W. Xu, W. Trappe, and Y. Zhang, "Anti-jamming timing channels for wireless networks," in *Proceedings of the first ACM conference on Wireless network security*, Alexandria, Virginia, pp. 203–213, 2008.

[10] S. D'Oro, L. Galluccio, G. Morabito, S. Palazzo, L. Chen, and F. Martignon, "Defeating Jamming With the Power of Silence: A Game-Theoretic Analysis," *IEEE Transactions on Wireless Communications*, vol. 14, no. 5, pp. 2337–2352, May 2015.

[11] M. Lichtman, J. D. Poston, S. Amuru, C. Shahriar, T. C. Clancy, R. M. Buehrer, and J. H. Reed, "A Communications Jamming Taxonomy," *IEEE Security and Privacy*, vol. 14, no. 1, pp. 47–54, Feb 2016.

[12] M. Wilhelm, I. Martinovic, J. B. Schmitt, and V. Lenders, "Short paper: Reactive jamming in wireless networks: How realistic is the threat?" in *Proceedings of the ACM WiSec*, Hamburg, Germany, 2011.

[13] W. Xu, K. Ma, W. Trappe, and Y. Zhang, "Jamming sensor networks: Attack and defense strategies," *IEEE Network*, vol. 20, no. 3, pp. 41–47, 2006.

[14] W. Xu, W. Trappe, Y. Zhang, and T. Wood, "The feasibility of launching and detecting jamming attacks in wireless networks," in *Proceedings of the 6th ACM International Symposium on Mobile Ad Hoc Networking and Computing*, Urbana-Champaign, IL, 2005.

[15] S. Amuru and R. M. Buehrer, "Optimal Jamming Strategies in Digital Communications–Impact of Modulation," in *IEEE Global Communications Conference*, Austin, TX, Dec 2014.

[16] S. Verdu and T. Han, "A general formula for channel capacity," *IEEE Transactions on Information Theory*, vol. 40, no. 4, pp. 1147–1157, 1994.

[17] D. J. Torrieri, "Fundamental limitations on repeater jamming of frequency-hopping communications," *IEEE Journal on Selected Areas in Communications*, vol. 7, no. 4, pp. 569–575, 1989.

[18] D. L. Adamy, *EW 102: A Second Course in Electronic Warfare*. Artech House, 2004.

[19] C. Ko, H. Nguyen-Le, and L. Huang, "ML-based follower jamming rejection in slow FH/MFSK systems with an antenna array," *IEEE Transactions on Communications*, vol. 56, no. 9, pp. 1536–1544, 2008.

[20] C. J. Burges, "A tutorial on support vector machines for pattern recognition," *Data Mining and Knowledge Discovery*, vol. 2, no. 2, pp. 121–167, 1998.

[21] R. Caruana and A. Niculescu-Mizil, "An empirical comparison of supervised learning algorithms," in *Proceedings of the 23rd International Conference on Machine Learning*, Pittsburgh, PA, 2006.

[22] M. Z. Win and J. H. Winters, "Analysis of hybrid selection/maximal-ratio combining of diversity branches with unequal SNR in Rayleigh fading," in *IEEE Vehicular Technology Conference*, Amsterdam, the Netherlands, 1999.

[23] C. E. Shannon and W. Weaver, "A Mathematical Theory of Communication," *University of Illinois Press*, 1949.

[24] P. Anghel, et al., "Exact symbol error probability of a cooperative network in a Rayleigh-fading environment," *IEEE Transactions on Wireless Communications*, vol. 3, no. 5, pp. 1416–1421, 2004.

[25] J. N. Laneman and G. W. Wornell, "Energy-efficient antenna sharing and relaying for wireless networks," in *IEEE Wireless Communications and Networking Conference*, Chicago, IL, 2000.

[26] J. M. Wozencraft and I. M. Jacobs, *Principles of Communication Engineering*. Wiley New York, 1965, vol. 32.

[27] P. Lee, "Computation of the bit error rate of coherent M-ary PSK with Gray code bit mapping," *IEEE Transactions on Communications*, vol. 34, pp. 488–491, 1986.

[28] H. V. Poor, *An Introduction to Signal Detection and Estimation*. 1em Berlin: Springer, 1994.

[29] J. Proakis and M. Salehi, *Digital Communications*, ser. McGraw-Hill higher education. McGraw-Hill Education, 2007.

List of Acronyms

AMC	Adaptive Modulation And Coding
ASK	Amplitude-Shift Keying
AWGN	Additive White Gaussian Noise
CF	Crest Factor
CW	Continuous Wave
DRFM	Digital RF Memory
FDM	Frequency Division Multiplexing
FHSS	Frequency-Hopping Spread Spectrum
FSK	Frequency-Shift Keying
GLONASS	Global Navigation Satellite System
GPS	Global Positioning System
IEEE	Institute of Electrical and Electronics Enginers
JSR	Jammer-To-Signal Ratio
LDPC	Low-Density Parity-Check
LPI/LPD	Low Probability Of Intercept And Detection
LTE	Long-Term Evolution
MAC	Media Access Control
MIMO	Multiple-Input And Multiple-Output
MRC	Maximal-Ratio Combining
OOK	On-Off Keying
OSI	Open Systems Interconnection
PAM	Pulse-Amplitude Modulation
PAPR	Peak-To-Average Power Ratio
PPM	Pulse Position Modulation
PSK	Phase-Shift Keying
QAM	Quadrature Amplitude Modulation
RF	An Radio Frequency
SC	Selection Combining
SNR	Signal-To-Noise Ratio
SVM	Support Vector Machine
TDMA	Time Division Multiple Access

DISTRIBUTION LIST

DTIC/OCP
8725 John J. Kingman Rd, Suite 0944
Ft Belvoir, VA 22060-6218   1 cy

AFRL/RVIL
Kirtland AFB, NM 87117-5776   2 cys

Official Record Copy
AFRL/RVSW/Khanh Pham   1 cy

(This page intentionally left blank)